

Nachname/  
Last name

Vorname/  
First name

Matrikelnr./  
Matriculation no

## Scheinklausur 03.04.2019

- Bitte tragen Sie zuerst auf dem Deckblatt Ihren Namen, Ihren Vornamen und Ihre Matrikelnummer ein. Tragen Sie dann auf den anderen Blättern (auch auf Konzeptblättern) Ihre Matrikelnummer ein.

*Please fill in your last name, your first name, and your matriculation number on this page and fill in your matriculation number on all other pages (including draft pages).*

- Die Prüfung besteht aus 25 Blättern: 1 Deckblatt, 19 Aufgabenblättern mit insgesamt 3 Aufgaben und 5 Blättern Man-Pages.

*The examination consists of 25 pages: 1 cover sheet, 19 sheets containing 3 assignments, and 5 sheets for man pages.*

- Es sind keinerlei Hilfsmittel erlaubt!

*No additional material is allowed.*

- Die Prüfung gilt als nicht bestanden, wenn Sie versuchen, aktiv oder passiv zu betrügen.

*You fail the examination if you try to cheat actively or passively.*

- Sie können auch die Rückseite der Aufgabenblätter für Ihre Antworten verwenden. Wenn Sie zusätzliches Konzeptpapier benötigen, verständigen Sie bitte die Klausuraufsicht.

*You can use the back side of the task sheets for your answers. If you need additional draft paper, please notify one of the supervisors.*

- Bitte machen Sie eindeutig klar, was Ihre endgültige Lösung zu den jeweiligen Teilaufgaben ist. Teilaufgaben mit widersprüchlichen Lösungen werden mit 0 Punkten bewertet.

*Make sure to clearly mark your final solution to each question. Questions with multiple, contradicting answers are void (0 points).*

- Programmieraufgaben sind gemäß der Vorlesung in C zu lösen.

*Programming assignments have to be solved in C.*

Die folgende Tabelle wird von uns ausgefüllt!

*The following table is completed by us!*

Aufgabe	1	2	3	Total
Max. Punkte	20	20	20	60
Erreichte Punkte				
Note				

## Aufgabe 1: C Grundlagen

*Assignment 1: C Basics*

a) Betrachten Sie die folgende Funktion `calculate()`.

*Consider the following function `calculate()`.*

```
1 double calculate(int *ptr, int n) {
2   assert(n != 0);
3   int result = 0;
4   int *ptr2 = ptr + n;
5   while (ptr < ptr2)
6     result += *ptr++;
7   return result / n;
8 }
```

Was berechnet die Funktion?

**1 pt**

*What does the function calculate?*

---

---

---

---

Welches Problem kann bei der Addition `+=` in Zeile 6 auftreten? Wie könnte dieses Problem zur Laufzeit erkannt werden?

**2 pt**

*Which issue might occur when the addition `+=` in line 6 is executed? How could this issue be detected at runtime?*

---

---

---

---

---

---

---

---

Welches Problem kann bei der Division in Zeile 7 auftreten? Wie kann das Problem vermieden werden?

**2 pt**

*Which issue might occur when executing the division in line 7? How could this issue be avoided?*

---

---

---

---

---

---

---

---

---

---

- b) Nennen Sie eine C-Bibliotheksfunktion, die nur manchmal einen Systemaufruf durchführt. Erklären Sie, in welchen Fällen ein Systemaufruf durchgeführt wird.

**1 pt**

*Name a C library function which does not always perform a system call. Explain the cases in which the function performs a system call.*

---

---

---

---

---

---

- c) Betrachten Sie die folgende Funktionen `a()` und `b()`.

*Consider the following functions `a()` und `b()`.*

```
void set_configuration(struct config *);
```

```
void a() {  
    struct config c;  
    memset(&c, 0, sizeof(struct config));  
    c.verbose = 1;  
    set_configuration(&c);  
}
```

```
void b() {  
    struct config *c = malloc(sizeof(struct config));  
    memset(c, 0, sizeof(struct config));  
    c->verbose = 1;  
    set_configuration(c);  
}
```

Begründen Sie, warum es nicht möglich ist, eine Funktion `set_configuration()` zu schreiben, die von beiden Funktionen `a()` und `b()` korrekt aufgerufen werden kann.

**2 pt**

*Explain why it is not possible to write a function `set_configuration()` that both `a()` and `b()` can call correctly.*



d) In dieser Aufgabe sollen Sie Funktionen zur Nutzung eines Bitvektors implementieren.

- Ein Bitvektor speichert eine Menge an Flags.
- Jedes Flag wird durch genau ein Bit dargestellt. Wenn das Flag gesetzt ist, hat das Bit den Wert 1, ansonsten 0.
- Die einzelnen Flags werden mit einem Index assoziiert.
- Ihre Bitvektor-Implementierung soll eine beliebige, aber nach der Initialisierung feste Anzahl an Flags speichern können.

*In this part, you will implement functions for using a bit vector.*

- *A bit vector saves a set of flags.*
- *Every flag is represented by exactly one bit in the bit vector. If the flag is set, the bit is set to 1, otherwise the bit is set to 0.*
- *Every flag is associated with an index.*
- *Your bit vector implementation should save an arbitrary number of flags that is fixed after initialization.*

Wählen Sie einen geeigneten Integer-Typ für den Bitvektor.

**0.5 pt**

- Der Typ sollte so gewählt werden, dass auf heute üblichen Prozessorarchitekturen effizient damit gerechnet werden kann.
- Achten Sie darauf, dass Ihre Implementierung der übrigen Teilaufgaben korrekt mit dem hier gewählten Typ funktioniert. Schreiben Sie portablen Code.

*Choose an appropriate integer type for the bit vector.*

- *Choose the type so that efficient operations are possible on current processor architectures.*
- *Make sure your implementation of the remaining parts of this exercise works correctly with the integer type you choose. Write portable code.*

**typedef**

```
.....  
.....  
    bv_int;  
.....
```

Vervollständigen Sie die Definition von `bpe`. Die Konstante soll die Anzahl an Bits in einem Element des Bitvektors enthalten.

**0.5 pt**

*Complete the definition of `bpe`. The constant should contain the number of bits per element of the bit vector.*

```
/* Number of bits per element */  
.....
```

```
static const size_t bpe =  
.....
```

















---

---

---

---

---

---

---

---

---

---

c) Vervollständigen Sie die Funktion `readFileList()`, welche die Textdatei mit der Liste der zu prüfenden Dateien einliest. Für jede aufgelistete Datei prüft die Funktion, ob die Datei seit der letzten Synchronisation mit dem Server verändert wurde.

**8 pt**

- Die Funktion erhält als Argument den absoluten Pfad zur einzulesenden Textdatei.
- Jede Zeile der Textdatei enthält den absoluten Pfad zu einer zu überprüfenden Datei.
- Wenn eine Datei nicht modifiziert wurde, sollen Sie den Pfad der Datei und “not modified” auf der Konsole ausgeben. Ansonsten sollen Sie den Pfad und “modified or unknown” ausgeben.
- Verwenden Sie `fgets()`, um die Textdatei zeilenweise einzulesen. Berücksichtigen Sie, dass `fgets()` den newline-Charakter ‘\n’ am Ende jeder Zeile nicht abschneidet, dieser jedoch nicht Teil des Pfades ist.
- Sie dürfen annehmen, dass jede Zeile der Textdatei maximal `MAX_PATH` Zeichen enthält.
- Leere Zeilen in der Textdatei sollen ignoriert werden.

*Complete the function `readFileList()` which reads the text file with the list of files to be checked. For each listed file, the function checks, whether the file was changed since the last synchronization with the server.*

- *The function is passed the absolute path to the text file to be read.*
- *Each line of the text file contains an absolute path to a file to be checked.*
- *If a file was not modified, you shall print the path of the file and “not modified” to the console. Otherwise the path and “modified or unknown” shall be printed.*
- *Use `fgets()` to read the text file line by line. Note that `fgets()` does not remove the newline character ‘\n’ at the end of each line. This newline character is not part of the file path.*
- *You may assume that each line of the text file contains at most `MAX_PATH` characters.*
- *Empty lines in the text file shall be ignored.*



## Aufgabe 3: Multilevel Feedback Queue

### Assignment 3: Multilevel Feedback Queue

Ein Multilevel Feedback Queue Scheduler (MLFB) bestimmt Prioritäten dynamisch basierend darauf, ob die Tasks ihre Zeitscheiben voll ausnutzen oder nicht. Implementieren Sie einen solchen Scheduler mit den folgenden Eigenschaften:

- Der Scheduler unterscheidet zwischen drei unterschiedlichen Prioritäten. Zwischen Task gleicher Priorität wird Round-Robin-Scheduling verwendet.
- Wenn ein Task seine Zeitscheibe vollständig ausnutzt, wird seine Priorität um eins reduziert.
- Wenn ein Task seine Zeitscheibe dreimal direkt in Folge nicht vollständig ausnutzt, wird seine Priorität um eins erhöht.
- Sie können davon ausgehen, dass zu jedem Zeitpunkt mindestens ein Task lauffähig ist, d.h. mindestens eine Queue ist nicht leer.
- Alle Zeiten werden als Vielfaches der Timer-Periode angegeben.

*A multilevel feedback queue scheduler (MLFB) determines priorities dynamically depending on whether tasks fully utilize their time slices or not. Implement such a scheduler with the following properties:*

- *The scheduler distinguishes between three different priorities. Round-robin scheduling is used for tasks of equal priority.*
- *If a task fully uses its time slice, its priority is reduced by one.*
- *If a task does not fully use its time slice three times directly in a row, its priority is increased by one.*
- *You can assume that at all times at least one task is runnable, i.e., at least one queue is non-empty.*
- *All times are stated as multiples of the timer period.*

```

struct queue_entry {
    struct queue_entry *next;
    struct queue_entry *prev;
};

struct task {
    uint32_t ts;           // Remaining part of current time slice
    uint32_t prio;        // Priority of the task (index into queue array)
    struct queue_entry rq; // Scheduler queue links

    uint32_t block_cnt; // Counter of partially unused timeslices - when the
                        // counter reaches 3, increase the priority of the task.
                        // The counter is set to 0 whenever the timeslice expires.

    // ... (task context omitted for brevity)
};

// list of runnable (not blocked) tasks of the same priority
struct runqueue {
    struct queue_entry head;
};

```


```
#define NUM_PRIIO 3 // Number of different priorities
                    // (2 is the highest priority, 0 the lowest)
// Length of the time slice of each priority
const int TS_LEN[NUM_PRIIO] = { 20, 10, 5 };

struct runqueue rqs[NUM_PRIIO]; // MLFB queues (one per priority)
struct task *current;           // Currently executed task


void context_switch(struct task *t); // Performs a context switch
                                     // if the task has changed.
```

a) Jede der Queues ist, wie in der folgenden Abbildung zu sehen, als doppelt verketete Liste realisiert. Leere Queues enthalten der Einfachheit halber einen Zyklus, sodass `next` und `prev` des Listenkopfes auf den Listenkopf selbst zeigen.

*Each of the queues is, as shown in the following figure, implemented as a double linked list. For the sake of simplicity, empty queues contain a cycle, so that `next` and `prev` of the list head point back to the list head.*



(a) Empty task queue



(b) Task queue with two tasks

Vervollständigen Sie die Funktion `from_task()`, die aus einem Pointer auf einen Task den Pointer auf den Queue-Eintrag (`rq`-Feld) des Tasks berechnet.

**1.0 pt**

*Complete the function `from_task()` which takes a pointer to a task and calculates the pointer to the queue entry (`rq` field) of the task.*

```
struct queue_entry *from_task(struct task *t) {
.....
.....
.....
.....
}
```

Vervollständigen Sie die Funktion `from_list_entry()`, die umgekehrt aus einem Pointer auf einen Queue-Eintrag den Pointer auf den dazugehörigen Task berechnet. Nehmen Sie dabei an, dass der Compiler die Reihenfolge der Felder der Struktur nicht ändert und kein Padding zwischen den Feldern einfügt.

**1.5 pt**

*Complete the function `from_list_entry()` which vice versa takes a pointer to a queue entry and computes the pointer to the corresponding task. Assume that the compiler does not reorder the fields of the struct and does not insert padding between the fields.*











**NAME**  
fclose – close a stream

**SYNOPSIS**  
#include <stdio.h>

int fclose(FILE \*fp);

**DESCRIPTION**

The **fclose()** function will flushes the stream pointed to by *fp* (writing any buffered output data using **flush(3)**) and closes the underlying file descriptor.

**RETURN VALUE**

Upon successful completion 0 is returned. Otherwise, **EOF** is returned and *errno* is set to indicate the error. In either case any further access (including another call to **fclose()**) to the stream results in undefined behavior.

**ERRORS**

**EBADF**

The file descriptor underlying *fp* is not valid.

The **fclose()** function may also fail and set *errno* for any of the errors specified for the routines **close(2)**, **write(2)** or **flush(3)**.

**NAME**

fgetc, fgets, getc, getchar, gets, ungetc – input of characters and strings

**SYNOPSIS**

#include <stdio.h>

char \*fgets(char \*s, int size, FILE \*stream);

**DESCRIPTION**

**fgets()** reads in at most one less than *size* characters from *stream* and stores them into the buffer pointed to by *s*. Reading stops after an **EOF** or a newline. If a newline is read, it is stored into the buffer. A '\0' is stored after the last character in the buffer.

**RETURN VALUE**

**fgets()** returns *s* on success, and NULL on error or when end of file occurs while no characters have been read.

**NAME**

fopen – stream open function

**SYNOPSIS**

#include <stdio.h>

FILE \*fopen(const char \*path, const char \*mode);

**DESCRIPTION**

The **fopen()** function opens the file whose name is the string pointed to by *path* and associates a stream with it.

The argument *mode* points to a string beginning with one of the following sequences (Additional characters may follow these sequences.):

**r** Open text file for reading. The stream is positioned at the beginning of the file.

**r+** Open for reading and writing. The stream is positioned at the beginning of the file.

**w** Truncate file to zero length or create text file for writing. The stream is positioned at the beginning of the file.

**w+** Open for reading and writing. The file is created if it does not exist, otherwise it is truncated. The stream is positioned at the beginning of the file.

The *mode* string can also include the letter 'b' either as a last character or as a character between the characters in any of the two-character strings described above. This is strictly for compatibility with C89 and has no effect; the 'b' is ignored on all POSIX conforming systems, including Linux. (Other systems may treat text files and binary files differently, and adding the 'b' may be a good idea if you do I/O to a binary file and expect that your program may be ported to non-UNIX environments.)

**RETURN VALUE**

Upon successful completion **fopen()** returns a *FILE* pointer. Otherwise, NULL is returned and *errno* is set to indicate the error.

**ERRORS**

**EINVAL**

The *mode* provided to **fopen()**, **fdopen()**, or **freopen()** was invalid.

The **fopen()** function may also fail and set *errno* for any of the errors specified for the routine **open(2)**.

MALLOCC(3)	Linux Programmer's Manual	MALLOCC(3)	
<b>NAME</b>	malloc, free — allocate and free dynamic memory		
<b>SYNOPSIS</b>	<pre>#include &lt;stdlib.h&gt; void *malloc(size_t size); void free(void *ptr);</pre>		
<b>DESCRIPTION</b>	<p>The <code>malloc()</code> function allocates <i>size</i> bytes and returns a pointer to the allocated memory. <i>The memory is not initialized.</i> If <i>size</i> is 0, then <code>malloc()</code> returns either NULL, or a unique pointer value that can later be successfully passed to <code>free()</code>.</p> <p>The <code>free()</code> function frees the memory space pointed to by <i>ptr</i>, which must have been returned by a previous call to <code>malloc()</code>. Otherwise, or if <i>free(ptr)</i> has already been called before, undefined behavior occurs. If <i>ptr</i> is NULL, no operation is performed.</p> <p><b>RETURN VALUE</b></p> <p>The <code>malloc()</code> function return a pointer to the allocated memory, which is suitably aligned for any built-in type. On error, these functions return NULL. NULL may also be returned by a successful call to <code>malloc()</code> with a <i>size</i> of zero.</p> <p>The <code>free()</code> function returns no value.</p>		
<b>ERRORS</b>	<code>malloc()</code> can fail with the following error:		
<b>ENOMEM</b>	Out of memory. Possibly, the application hit the <code>RLIMIT_AS</code> or <code>RLIMIT_DATA</code> limit described in <code>getrlimit(2)</code> .		
<b>NOTES</b>	<p>By default, Linux follows an optimistic memory allocation strategy. This means that when <code>malloc()</code> returns non-NULL there is no guarantee that the memory really is available. In case it turns out that the system is out of memory, one or more processes will be killed by the OOM killer. For more information, see the description of <i>/proc/sys/vm/overcommit_memory</i> and <i>/proc/sys/vm/oom_adj</i> in <code>proc(5)</code>, and the Linux kernel source file <i>Documentation/vm/overcommit-accounting</i>.</p> <p>Normally, <code>malloc()</code> allocates memory from the heap, and adjusts the size of the heap as required, using <code>sbrk(2)</code>. When allocating blocks of memory larger than <code>MMAP_THRESHOLD</code> bytes, the glibc <code>malloc()</code> implementation allocates the memory as a private anonymous mapping using <code>mmap(2)</code>. <code>MMAP_THRESHOLD</code> is 128 kB by default, but is adjustable using <code>mallopt(3)</code>. Prior to Linux 4.7 allocations performed using <code>mmap(2)</code> were unaffected by the <code>RLIMIT_DATA</code> resource limit; since Linux 4.7, this limit is also enforced for allocations performed using <code>mmap(2)</code>.</p> <p>To avoid corruption in multithreaded applications, mutexes are used internally to protect the memory-management data structures employed by these functions. In a multithreaded application in which threads simultaneously allocate and free memory, there could be contention for these mutexes. To scalably handle memory allocation in multithreaded applications, glibc creates additional <i>memory allocation arenas</i> if mutex contention is detected. Each arena is a large region of memory that is internally allocated by the system (using <code>brk(2)</code> or <code>mmap(2)</code>), and managed with its own mutexes.</p> <p>SUSv2 requires <code>malloc()</code> to set <code>errno</code> to <code>ENOMEM</code> upon failure. Glibc assumes that this is done (and the glibc versions of these routines do this); if you use a private malloc implementation that does not set <code>errno</code>, then certain library routines may fail without having a reason in <code>errno</code>.</p> <p>Crashes in <code>malloc()</code> or <code>free()</code> are almost always related to heap corruption, such as overflowing an allocated chunk or freeing the same pointer twice.</p> <p>The <code>malloc()</code> implementation is tunable via environment variables; see <code>mallopt(3)</code> for details.</p>		
MALLOCC(3)	Linux Programmer's Manual		
<b>NAME</b>	memset — fill memory with a constant byte		
<b>SYNOPSIS</b>	<pre>#include &lt;string.h&gt; void *memset(void *s, int c, size_t n);</pre>		
<b>DESCRIPTION</b>	The <code>memset()</code> function fills the first <i>n</i> bytes of the memory area pointed to by <i>s</i> with the constant byte <i>c</i> .		
<b>RETURN VALUE</b>	The <code>memset()</code> function returns a pointer to the memory area <i>s</i> .		
<b>NAME</b>	memcpy — copy memory area		
<b>SYNOPSIS</b>	<pre>#include &lt;string.h&gt; void *memcpy(void *dest, const void *src, size_t n);</pre>		
<b>DESCRIPTION</b>	The <code>memcpy()</code> function copies <i>n</i> bytes from memory area <i>src</i> to memory area <i>dest</i> . The memory areas must not overlap. Use <code>memmove(3)</code> if the memory areas do overlap.		
<b>RETURN VALUE</b>	The <code>memcpy()</code> function returns a pointer to <i>dest</i> .		
<b>NOTES</b>	Failure to observe the requirement that the memory areas do not overlap has been the source of significant bugs. (POSIX and the C standards are explicit that employing <code>memcpy()</code> with overlapping areas produces undefined behavior.) Most notably, in glibc 2.13 a performance optimization of <code>memcpy()</code> on some platforms (including x86-64) included changing the order in which bytes were copied from <i>src</i> to <i>dest</i> .		
<b>NAME</b>	strlen — calculate the length of a string		
<b>SYNOPSIS</b>	<pre>#include &lt;string.h&gt; size_t strlen(const char *s);</pre>		
<b>DESCRIPTION</b>	The <code>strlen()</code> function calculates the length of the string pointed to by <i>s</i> , excluding the terminating null byte ( <code>\0</code> ).		
<b>RETURN VALUE</b>	The <code>strlen()</code> function returns the number of characters in the string pointed to by <i>s</i> .		

**NAME**  
mq\_close – close a message queue descriptor

**SYNOPSIS**  
#include <mqqueue.h>

int mq\_close(mqd\_t mqdes);

Link with `-lrt`.

**DESCRIPTION**  
mq\_close() closes the message queue descriptor *mqdes*.

If the calling process has attached a notification request to this message queue via *mqdes*, then this request is removed, and another process can now attach a notification request.

**RETURN VALUE**  
On success **mq\_close()** returns 0; on error, `-1` is returned, with *errno* set to indicate the error.

**ERRORS**  
**EBADF**  
The descriptor specified in *mqdes* is invalid.

**NOTES**  
All open message queues are automatically closed on process termination, or upon **execve(2)**.

**NAME**  
mq\_open – open a message queue

**SYNOPSIS**  
#include <fcntl.h> /\* For O\_\* constants \*/  
#include <mqqueue.h>

mqd\_t mq\_open(const char \*name, int oflag);

**DESCRIPTION**  
mq\_open() creates a new POSIX message queue or opens an existing queue. The queue is identified by *name*.

The *oflag* argument specifies flags that control the operation of the call. (Definitions of the flags values can be obtained by including <fcntl.h>.) Exactly one of the following must be specified in *oflag*:

**O\_RDONLY**  
Open the queue to receive messages only.

**O\_WRONLY**  
Open the queue to send messages only.

**O\_RDWR**  
Open the queue to both send and receive messages.

Zero or more of the following flags can additionally be *OR*ed in *oflag*:

**O\_NONBLOCK**  
Open the queue in nonblocking mode. In circumstances where **mq\_receive(3)** and **mq\_send(3)** would normally block, these functions instead fail with the error **EAGAIN**.

**RETURN VALUE**  
On success, **mq\_open()** returns a message queue descriptor for use by other message queue functions. On error, **mq\_open()** returns (*mqd\_t*) `-1`, with *errno* set to indicate the error.

**ERRORS**  
**EACCES**  
The queue exists, but the caller does not have permission to open it in the specified mode.  
**ENOENT**  
The **O\_CREAT** flag was not specified in *oflag*, and no queue with this *name* exists. *mode*.

**NAME**

mq\_send – send a message to a message queue

**SYNOPSIS**

```
#include <mqqueue.h>
```

```
int mq_send(mqd_t mqdes, const char *msg_ptr,
            size_t msg_len, unsigned int msg_prio);
```

**DESCRIPTION**

**mq\_send()** adds the message pointed to by *msg\_ptr* to the message queue referred to by the descriptor *mqdes*. The *msg\_len* argument specifies the length of the message pointed to by *msg\_ptr*; this length must be less than or equal to the queue's *mq\_msgsize* attribute. Zero-length messages are allowed.

The *msg\_prio* argument is a nonnegative integer that specifies the priority of this message. Messages are placed on the queue in decreasing order of priority, with newer messages of the same priority being placed after older messages with the same priority.

If the message queue is already full (i.e., the number of messages on the queue equals the queue's *mq\_maxmsg* attribute), then, by default, **mq\_send()** blocks until sufficient space becomes available to allow the message to be queued, or until the call is interrupted by a signal handler. If the **O\_NONBLOCK** flag is enabled for the message queue description, then the call instead fails immediately with the error **EAGAIN**.

**RETURN VALUE**

On success, **mq\_send()** returns zero; on error,  $-1$  is returned, with *errno* set to indicate the error.

**ERRORS****EAGAIN**

The queue was full, and the **O\_NONBLOCK** flag was set for the message queue description referred to by *mqdes*.

**EBADF**

The descriptor specified in *mqdes* was invalid.

**EINTR**

The call was interrupted by a signal handler; see **signal(7)**.

**EMSGSIZE**

*msg\_len* was greater than the *mq\_msgsize* attribute of the message queue.

**NAME**

mq\_receive – receive a message from a message queue

**SYNOPSIS**

```
#include <mqqueue.h>
```

```
ssize_t mq_receive(mqd_t mqdes, char *msg_ptr,
                  size_t msg_len, unsigned int *msg_prio);
```

**DESCRIPTION**

**mq\_receive()** removes the oldest message with the highest priority from the message queue referred to by the descriptor *mqdes*, and places it in the buffer pointed to by *msg\_ptr*. The *msg\_len* argument specifies the size of the buffer pointed to by *msg\_ptr*; this must be greater than or equal to the *mq\_msgsize* attribute of the queue (see **mq\_getattr(3)**). If *msg\_prio* is not NULL, then the buffer to which it points is used to return the priority associated with the received message.

If the queue is empty, then, by default, **mq\_receive()** blocks until a message becomes available, or the call is interrupted by a signal handler. If the **O\_NONBLOCK** flag is enabled for the message queue description, then the call instead fails immediately with the error **EAGAIN**.

**RETURN VALUE**

On success, **mq\_receive()** returns the number of bytes in the received message; on error,  $-1$  is returned, with *errno* set to indicate the error.

**ERRORS****EAGAIN**

The queue was empty, and the **O\_NONBLOCK** flag was set for the message queue description referred to by *mqdes*.

**EBADF**

The descriptor specified in *mqdes* was invalid.

**EINTR**

The call was interrupted by a signal handler; see **signal(7)**.

**EMSGSIZE**

*msg\_len* was less than the *mq\_msgsize* attribute of the message queue.



**NAME** printf, fprintf, sprintf, snprintf, vprintf, vsprintf, vsnprintf – formatted output conversion

**SYNOPSIS**

```
#include <stdio.h>
```

```
int printf(const char *format, ...);
int fprintf(FILE *stream, const char *format, ...);
```

**DESCRIPTION**

The functions in the **printf()** family produce output according to a *format* as described below. The function **printf()** writes output to *stdout*, the standard output stream; **fprintf()** writes output to the given output *stream*.

These functions write the output under the control of a *format* string that specifies how subsequent arguments are converted for output.

**Return value**

Upon successful return, these functions return the number of characters printed (not including the trailing '\0' used to end output to strings).

If an output error is encountered, a negative value is returned.

**Format of the format string**

The format string is a character string, beginning and ending in its initial shift state, if any. The format string is composed of zero or more directives: ordinary characters (not %), which are copied unchanged to the output stream; and conversion specifications, each of which results in fetching zero or more subsequent arguments. Each conversion specification is introduced by the character %, and ends with a *conversion specifier*.

The arguments must correspond properly (after type promotion) with the conversion specifier. By default, the arguments are used in the order given, where each conversion specifier asks for the next argument (and it is an error if insufficiently many arguments are given).

**The conversion specifier**

A character that specifies the type of conversion to be applied. The conversion specifiers and their meanings are:

- d, i** The *int* argument is converted to signed decimal notation.
- c** The *int* argument is converted to an *unsigned char*, and the resulting character is written.
- s** The *const char\** argument is expected to be a pointer to an array of character type (pointer to a string). Characters from the array are written up to (but not including) a terminating null byte ('\0').

**NAME** stat, lstat – get file status

**SYNOPSIS**

```
#include <sys/types.h>
#include <sys/stat.h>
#include <unistd.h>
```

```
int stat(const char *pathname, struct stat *buf);
int lstat(int fd, struct stat *buf);
```

**DESCRIPTION**

These functions return information about a file, in the buffer pointed to by *buf*. No permissions are required on the file itself, but—in the case of **stat()**—execute (search) permission is required on all of the directories in *pathname* that lead to the file.

**stat()** retrieves information about the file pointed to by *pathname*.

**stat()** is identical to **lstat()**, except that the file about which information is to be retrieved is specified by the file descriptor *fd*.

All of these system calls return a *stat* structure, which contains the following fields:

```
struct stat {
    dev_t  st_dev;      /* ID of device containing file */
    ino_t  st_ino;     /* inode number */
    mode_t st_mode;    /* file type and mode */
    nlink_t st_nlink;  /* number of hard links */
    uid_t  st_uid;     /* user ID of owner */
    gid_t  st_gid;     /* group ID of owner */
    dev_t  st_rdev;    /* device ID (if special file) */
    off_t  st_size;    /* total size, in bytes */
    blksize_t st_blksize; /* blocksizes for filesystem I/O */
    blkcnt_t st_blocks; /* number of 512B blocks allocated */

    time_t st_atime;   /* Time of last access */
    time_t st_mtime;   /* Time of last modification */
    time_t st_ctime;   /* Time of last status change */
};
```

The field *st\_atime* is changed by file accesses, for example, by **execve(2)**, **mknod(2)**, **pipe(2)**, **utime(2)**, and **read(2)** (of more than zero bytes). Other routines, like **mmap(2)**, may or may not update *st\_atime*.

The field *st\_mtime* is changed by file modifications, for example, by **mknod(2)**, **truncate(2)**, **utime(2)** and **write(2)** (of more than zero bytes). Moreover, *st\_mtime* of a directory is changed by the creation or deletion of files in that directory. The *st\_mtime* field is *not* changed for changes in owner, group, hard link count, or mode.

The field *st\_ctime* is changed by writing or by setting inode information (i.e., owner, group, link count, mode, etc.).

**RETURN VALUE**

On success, zero is returned. On error, **-1** is returned, and *errno* is set appropriately.